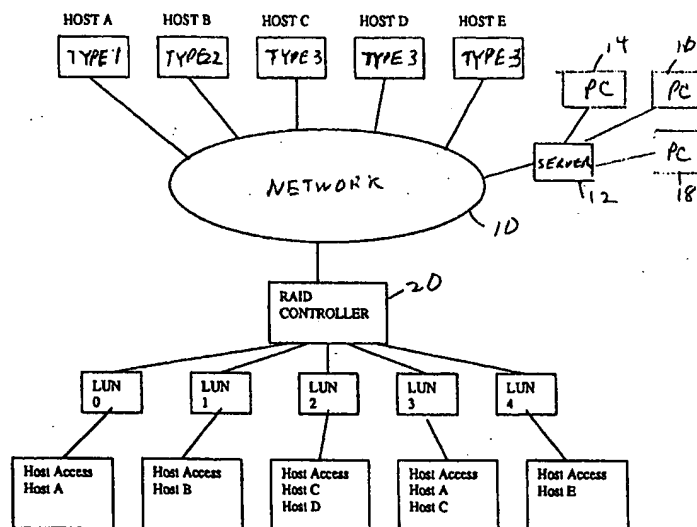




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 12/00, 12/14	A1	(11) International Publication Number: WO 00/29954 (43) International Publication Date: 25 May 2000 (25.05.00)
(21) International Application Number: PCT/US99/26450 (22) International Filing Date: 8 November 1999 (08.11.99) (30) Priority Data: 60/108,461 14 November 1998 (14.11.98) US (71) Applicant: MTI TECHNOLOGY CORPORATION [US/US]; 4905 East La Palma Avenue, Anaheim, CA 92807 (US). (72) Inventors: ROBBINS, Robert; 998-2 Alpine Terrace, Sunnyvale, CA 94086 (US). MEYER, Richard; 69 Devonshire Way, San Francisco, CA 94131 (US). SMITH, Jerry; 3643 Esther Street, Long Beach, CA 90804 (US). (74) Agents: HAUGHEY, Paul, C. et al.; Townsend And Townsend And Crew LLP, 8th Floor, Two Embarcadero Center, San Francisco, CA 94111-3834 (US).	(81) Designated States: CA, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: LOGICAL UNIT MAPPING IN A STORAGE AREA NETWORK (SAN) ENVIRONMENT



(57) Abstract

A system with a RAID controller (20) connected to a plurality of hosts (A-E), and having a storage area configurable as a plurality of logical units (0-4). The RAID controller (20) stores in an associated memory a table of access right for the hosts (A-E) to the logical units (0-4). Updates can be provided and additional material can be provided in new memory space that is subsequently made available to the various hosts (A-E) by providing updated access rights that include the new space. In one embodiment, the present invention also maps the logical unit number (0-4) from that requested by the host (A-E) to an actual number used in the RAID subsystem.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

provided in new memory space that is subsequently made available to the various hosts by providing updated access rights that include the new space.

In one embodiment, the present invention also maps the logical unit number from that requested by the host to an actual number used in the RAID subsystem. For example, multiple hosts may request logical unit number zero (LUN0), and other than the first host, they may be mapped to a different logical unit, such as LUN5, LUN8, etc.

For a further understanding of the nature and advantages of the invention, reference should be made to the following description taken in conjunction with the accompanying drawings.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a SAN network including a RAID controller according to the invention.

Fig. 2 is a flowchart illustrating the checking of access rights.

15

Fig. 3 is a block diagram illustrating two redundant RAID controllers.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

The invention allows the RAID controller user to allocate access to the storage area at the RAID subsystem level. Thus, a PC-NT host which will want to take control of everything it is connected to will only know about a designated portion of the memory space. The invention is particularly advantageous for a fast fibre optic network. In particular, previous network systems typically had a dedicated server that would control access to the memory subsystem. With faster links, this limits the access speed, and the present invention allows the elimination of this intermediate server.

25

Fig. 1 shows a network 10 connected to a number of host computers A-E. Also shown is a separate server 12, which may itself be connected to a number of separate hosts 14, 16, and 18. A RAID controller 20 is also connected to the network.

RAID controller 20 is shown as accessing multiple logical unit numbers 0-4. Illustrated below each logical unit number is an example of which hosts are granted access to that LUN.

30

The network in Fig. 1 is preferably a fibre network, but could be any other type of network or any switching or shared bus. For example, it could be a SCSI bus. The RAID controller of Fig. 1 is shown at a logical level, and could be two redundant

LOGICAL UNIT MAPPING IN A STORAGE AREA NETWORK (SAN) ENVIRONMENT

CROSS-REFERENCES TO RELATED APPLICATIONS

5 This application claims priority from U.S. provisional application no. 60/108,461 filed on November 14, 1998 and having the same title. The contents of this provisional application are incorporated herein by reference.

BACKGROUND OF THE INVENTION

10 The present invention relates to RAID (redundant array of independent disks) systems in a storage area network (SAN).

 In one typical RAID configuration, a RAID controller is connected to an array of disk drives. The other side of the RAID controller connects to multiple hosts through an intermediate server. In a SAN environment, the intermediate server is
15 eliminated, and the RAID controller is itself on the SAN network.

 In such a network, a PC-NT host would typically try to take control of all memory spaces to which it is connected. In addition, if the RAID memory was to be used as the boot memory, multiple PCs would want to use the same logical unit (LUN) address for booting. It would be desirable to have a RAID controller which can connect into an
20 SAN network and handle such operations transparent to the PCs on the network.

SUMMARY OF THE INVENTION

 The present invention provides a system with a RAID controller connected to a plurality of homogenous or heterogeneous hosts, and having a storage area
25 configurable as a plurality of logical units. The RAID controller stores in an associated memory a table of access rights for the hosts to the logical units.

 Preferably, a default setting is that a host does not have any access to the logical units of the RAID subsystem. An operator can then program the RAID controller to grant the desired amount of access to a particular host. Logical unit areas can be
30 reserved to add new hosts, for instance. In addition, new space can be provided to existing hosts dynamically without affecting operations in logical unit space already accessible by the particular host. Updates can be provided and additional material can be

RAID controllers as shown in Fig. 3. Each RAID controller can be connected to two channels to the host, channels 0 and 1.

RAID controllers must be able to supply common shared storage in both a heterogeneous and homogeneous networked environment without additional software required on the host systems to manage access to the units. The RAID controller must be able to partition the storage into logical units that can be presented to the hosts. This logical unit is a grouping of memory storage for which there is some method for mapping the access known by the RAID controller. It is independent of the actual physical storage media being used. For the RAID controller to provide the common storage in a SAN environment, it must be capable of limiting access to the logical units on a host by host basis.

The RAID controller software provides a method for controlling the access of the logical units on a host by host basis. It provides the user the capability of managing the logical units and limiting access. The user enters in the unique name or identifier for a specific host that can be obtained by the RAID controller on the network. The user then selects which logical units that host can be allowed to access. This information is maintained in a table in the controller's memory when the code is running, allowing the changes to be implemented dynamically. It is also saved in some form of memory storage (flash, EEPROM, NVRAM, disk, etc.) so that it is persistent across reboots, power failures, etc. A copy is also sent to all of the controller's failover partners in a redundant configuration to maintain access across failures.

The controller software allows the logical units to be accessible to any combination of hosts. The user can enable access by none, one, or any number of hosts as required. Because the controller manages the storage and access to it, the hosts on the SAN do not need to communicate with each other, or even know other hosts are on the network.

In addition to controlling whether or not the logical unit is accessible by a host, the user can also limit the host access to the logical unit to read only. This allows the user to set up the logical unit as shared by several selected hosts and specify which of those hosts can update the data and which can only view the data.

When the controller receives a request from a host for a logical unit, it checks the access table to determine if the host is allowed access. If the host is not allowed access, the controller reports back to the host that the logical unit is not connected. If the host is allowed access, the request is processed.

The user can access the RAID controller for configuration, etc., through a serial port connected to a terminal (console) or PC. Alternately, a PC connected through an Ethernet network could be used, or any other access method.

The flowchart of Fig. 2 illustrates how a host request 30 for access to the storage space is handled. The RAID controller first extracts from the host request the logical unit the host is trying to access and the host ID (step 32). The logical unit requested is compared to a table to determine if the host will be granted access rights (step 34).

If no access rights are provided, a response is sent to the host request indicating that the logical unit requested is not connected (step 36). If access is allowed, the request is serviced (step 38). A response is sent back to the host in either case (step 40).

The invention provides a high degree of configuration flexibility. For example, a user can configure the system to tell a particular host that there is not any logical unit available, fooling the host into thinking it is simply not connected. If multiple hosts are sharing a common storage, they can be granted the same access rights. The storage can be added to in an area where no access is provided, and then after the additions have been checked, access rights to the new storage can be dynamically added. Similarly, a portion can be dynamically removed from the access rights, modified, then access can be provided again. In addition to dynamically bringing storage in and out, hosts can be dynamically brought in and out by eliminating and adding access rights dynamically.

SAN volume mapping to specific host systems can be used to isolate storage to specific hosts systems for performance, management, and data protection.

Zones or host mappings support functions and business needs similar to how a traditional network can be subdivided into sub-nets. Performance, data protection, security, fire-walling, and different application requirements will drive zoning and host to LUN mapping.

Zoning can be accomplished via switches as well as via RAID devices using either port to port mappings, or LUN to World Wide Port Name mapping.

The present invention thus provides the ability to dynamically modify LUN access on the fly. Typically, this may be done to add access and expand storage without rebooting. Where a mirrored set is provided, both are written to.

The present invention can also provide limited access to a LUN, as opposed to total or no access. For example, read only access could be provided. This would be useful, for instance, to allow another host to do backup at the same time that a LUN is also accessible by other hosts. The backup server would be given read only access, and thus there is less chance of conflicts with other hosts attempting to read and write the same area. By granting a backup server read only access for backing up, it need only be given the access during the backup time and thus the backup server doesn't have to have access to the same memory all the time. This eliminates the need to use one of the hosts using the storage for backup purposes.

10 In addition, where a RAID server is used for an internet company, read only access could be provided. Another portion of the memory could be reserved for writing updates, which could then be subsequently provided to users by updating their access rights.

A customer can communicate with the RAID controller through a console connected to it or through a GUI interface from host. Both are preferably password protected.

The 6300 and 6700 both connect to hosts over a fibre channel interface. The 6300 described below is an example RAID controller using SCSI interfaces to a RAID array of disk drives that form the LUNs. The user manuals for the 6300 and 6700 describe the operation in more detail, and are available from MTI Technology Corp. The 6300 has a default condition where all hosts have total access. The 6700 is a RAID controller using fibre interfaces to the RAID disk drives. The 6700 has a default condition where no host has access. Any default condition could be used for any system configuration.

25 The port names set forth conform to IEEE standard 802-1990 for naming.

Configuring LUN Affinity

The RAID controllers support defining which hosts on the Fibre Channel Arbitrated Loop have access to which LUNs. Each LUN can be set up to be accessed by any number of hosts. The hosts are defined by their worldwide unique port name. On the 6700 no host has access to any LUN unless there is an entry in the table. On the 6300 all hosts have access to all LUNs if no table exists. Currently this access must be set up through the serial console port. On the 6700 the changes will have effect the next time the host attempts to access a LUN that has been changed.

6700 Host Access Table

1. At the serial console port, type gladMenu to the main menu.
2. Select 1 to get to the Field Service menu.
- 5 3. Select 1 to get to the Configuration menu.
4. Select 6 to display the Host Access Table.
5. Select 7 to edit the Host Access Table.
6. Enter P to add or delete a host, or L to edit the LUN access for a host already in the table, at the prompt. If you select L, go to step 9.
- 10 7. If you selected P, enter D to delete a host, or A to add a host, at the prompt.
8. Select either channel 0 or 1 at the prompt, based upon which port the LUN is assigned to.
9. Enter the 8 byte worldwide unique name of the host.
- 15 10. Enter a 12 byte text description to identify the host.
11. For each LUN displayed as disabled that you want to enable type t.

Sample Host Access Entry Display:

Enter channel (0 or 1): 0

20 Hosts currently logged in:

20:00:00:e0:8b:00:72:03

Port name: 20:00:00:e0:8b:00:72:03

Host name: System One

LUN 0 is enabled

25 LUN 1 is disabled

LUN 2 is enabled

LUN 3 is disabled

LUN 4 is disabled

30 6300 Host Access Table

1. At the serial console port, type gladMenu to get to the main menu.
2. Select 1 to get to the Field Service menu.
3. Select 7 to get to the FC LUN MAP menu.

4. If you need the port information for the nodes on the FCAL, select 5 to display the logged-in hosts.
5. Select 1 to display the LUN map.
6. Select 2 to remove a host from the LUN map.
- 5 7. Select 3 to add a host to the LUN map.
8. After making changes, select 4 to save LUN map.
9. Select 6 to remove the LUN map to start from scratch. WARNING: All the data in the LUN MAP file will be destroyed.
10. Select 7 to get a copy of partner board's LUN MAP onto this board.

10

Sample Logged In Hosts Display:

Host Map

ALPA Node Name

1 100000a0f1000132

15 2 100000b2a0001094

ef 10000098bf221032

Sample LUN MAP Display:

LUN Map

ALPA	Host LUN	RAID LUN
1	0	0
1	1	1
1	2	2
1	3	3
25 2	0	0
ef	0	0
ef	1	4
ef	2	5
ef	3	6

30

RAID LUN is the actual LUN number assigned in the 6300 array.

Host LUN is the LUN number that is presented to the host for that RAID LUN.

As will be understood by those of skill in the art, the present invention may be embodied in other specific forms without departing from the essential characteristics thereof. Accordingly, the foregoing description is intended to be illustrative, but not limiting, of the scope of the invention which is set forth in the following claims.

WHAT IS CLAIMED IS:

- 1 1. A memory storage system comprising:
2 a RAID controller configured to connect to a plurality of hosts;
3 a storage area configurable as a plurality of logical units; and
4 a memory, connected to said RAID controller, for storing access rights of
5 said hosts to said logical units.
- 1 2. The system of claim 1 wherein said RAID controller is configured to
2 dynamically change said access rights in said memory.
- 1 3. The system of claim 1 wherein said memory stores default access rights
2 granting no access.
- 1 4. The system of claim 1 further comprising a fibre channel network
2 connecting said memory storage system to said plurality of hosts.
- 1 5. The system of claim 1 wherein said memory further stores a remapping
2 of logical units, so that multiple hosts can access said memory storage system specifying
3 an identical logical unit, and be directed to different logical units.
- 1 6. The system of claim 1 wherein a grouping of memory storage
2 corresponding to logical units assigned to a particular host is transparent to said particular
3 host.
- 1 7. The system of claim 1 further comprising a table in said memory, said
2 table storing both said access rights and a mapping to translated logical units.
- 1 8. A memory storage system comprising:
2 a RAID controller configured to connect to a plurality of hosts;
3 a storage area configurable as a plurality of logical units; and
4 a memory, connected to said RAID controller, for storing access rights of
5 said hosts to said logical units;
6 said RAID controller being configured to dynamically change said access
7 rights in said memory;
8 said memory storing default access rights granting no access; and

9 said memory further storing a remapping of logical units, so that multiple
10 hosts can access said memory storage system specifying an identical logical unit, and be
11 directed to different logical units.

1 9. The system of claim 8 further comprising a fibre channel network
2 connecting said memory storage system to said plurality of hosts.

1 10. A method for operating a memory storage system comprising:
2 connecting a RAID controller to a plurality of hosts;
3 said RAID controller being connected to a storage area configurable as a
4 plurality of logical units; and
5 storing access rights of said hosts to said logical units.

1 11. The method of claim 10 further comprising dynamically altering said
2 access rights.

1 12. The method of claim 11 wherein said dynamically altering comprises
2 adding access rights to a particular host without affecting existing access rights of said
3 particular host.

1 13. The method of claim 10 further comprising storing a default of no
2 access rights prior to granting access rights to a host.

1 14. The method of claim 10 further comprising reserving a designated
2 plurality of logical units for future hosts to be added, and not granting existing hosts
3 access to said designated plurality of logical units.

1 15. A method for operating a memory storage system comprising:
2 connecting a RAID controller to a plurality of hosts;
3 said RAID controller being connected to a storage area configurable as a
4 plurality of logical units;
5 storing access rights of said hosts to said logical units;
6 dynamically altering said access rights, wherein said dynamically altering
7 comprises adding access rights to a particular host without affecting existing access rights
8 of said particular host; and

9 storing a default of no access rights prior to granting access rights to a
10 host.

1 16. The method of claim 15 further comprising reserving a designated
2 plurality of logical units for future hosts to be added, and not granting existing hosts
3 access to said designated plurality of logical units.

MTI Company Confidential...

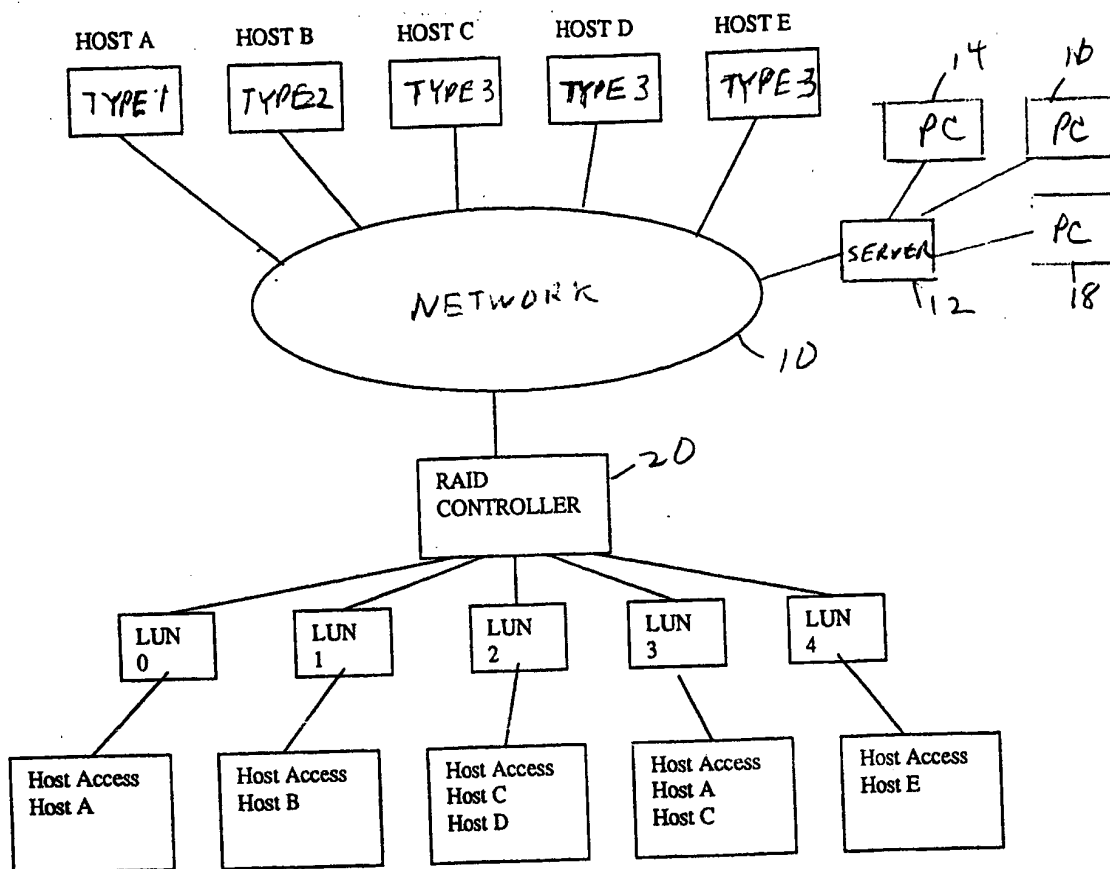


FIG. 1

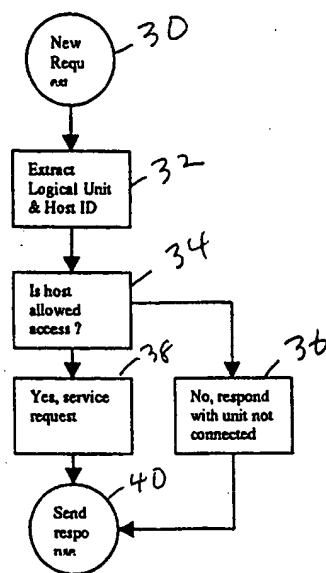


FIG. 2




FIG. 3

BEST AVAILABLE COPY

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US99/26450

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 12/00, 12/14

US CL : 711/114,164

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 711/114,164, 163

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WEST TEXT SEARCH, IEL ONLINE

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,552,776 A (WADE et al) 03 September 1996, col. 2, lines 32-58; col. 3, lines 5-55; col. 7, lines 19-59; col. 9-10.	1-16
A	US 5,812,754 A (LUI et al) 22 September 1998, abstract; fig. 1-2; col. 8, lines 9-34.	1-16
A	US 5,771,291 (NEWTON et al) 23 June 1998, col. 4, lines 11-26; col. 5, lines 8-35.	1-16

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 23 FEBRUARY 2000	Date of mailing of the international search report 17 MAR 2000
---	---

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JOHN W. CABEC *James R. Matthews*

Telephone No. (703) 305-0134

THIS PAGE BLANK (USPTO)